Choosing A Metric For Scanning Negative

A Presentation for the Interchange Format subcommittee of the Advanced Technology Committee of the Academy of Motion Picture Arts & Sciences Science and Technology Council prepared by Richard Patterson with a generous assist from Ed Giorgianni

A scanner functions in much the same way as a densitometer. Light passing through the film is detected by a sensor, and the variations in the amount of light are recorded as some form of data. The native metric for a scanner is some form of density-related measurement.

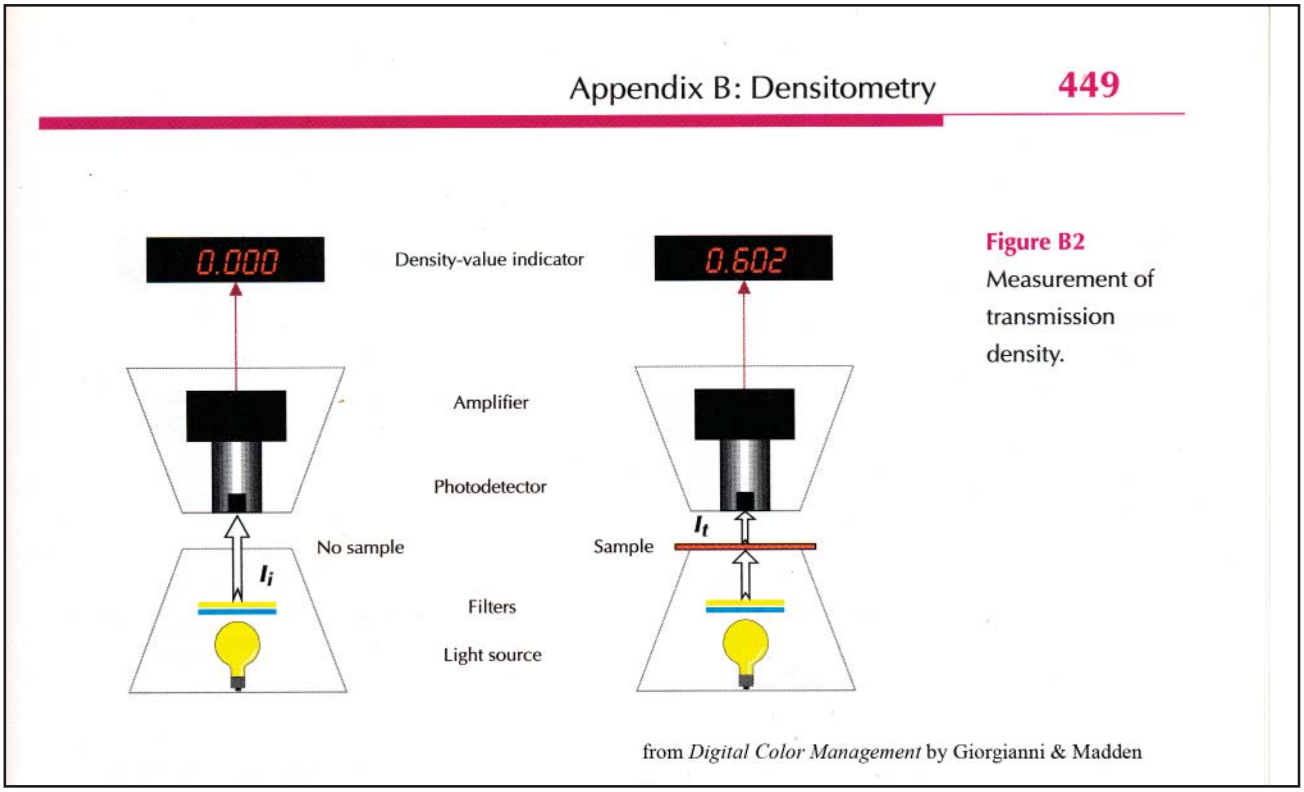

Optical Density is defined as the log10 of Opacity. Opacity is in turn defined as the reciprocal of Transmittance, and Transmittance is the "ratio of the transmitted radiance or luminous flux to the incident flux under specified conditions of irradiation." In other words Transmittance is the percentage or fraction of light that makes it through the film.

Because the perception of light, like the perception of sound, is largely a matter of the perception of ratios rather than absolute levels, it is more intuitive to use the logarithmic scale representing density rather than the linear scale representing opacity.

Transmittance	Opacity	Density
0.1	10.0	1.0000
0.2	5.0	0.6990
0.3	3.3	0.5229
0.4	2.5	0.3979
0.5	2.0	0.3010
0.6	1.7	0.2218
0.7	1.4	0.1549
0.8	1.3	0.0969
0.9	1.1	0.0458
1.0	1.0	0.0000

Like most substances Dyes used in films absorb light differently at different wavelengths. That is after all what gives the dye its color. (A sensitometry strip made by coating carbon black on a clear base absorbs the light equally at all wavelengths.) The way in which a dye absorbs different wavelengths can be represented by a spectral dye density curve.

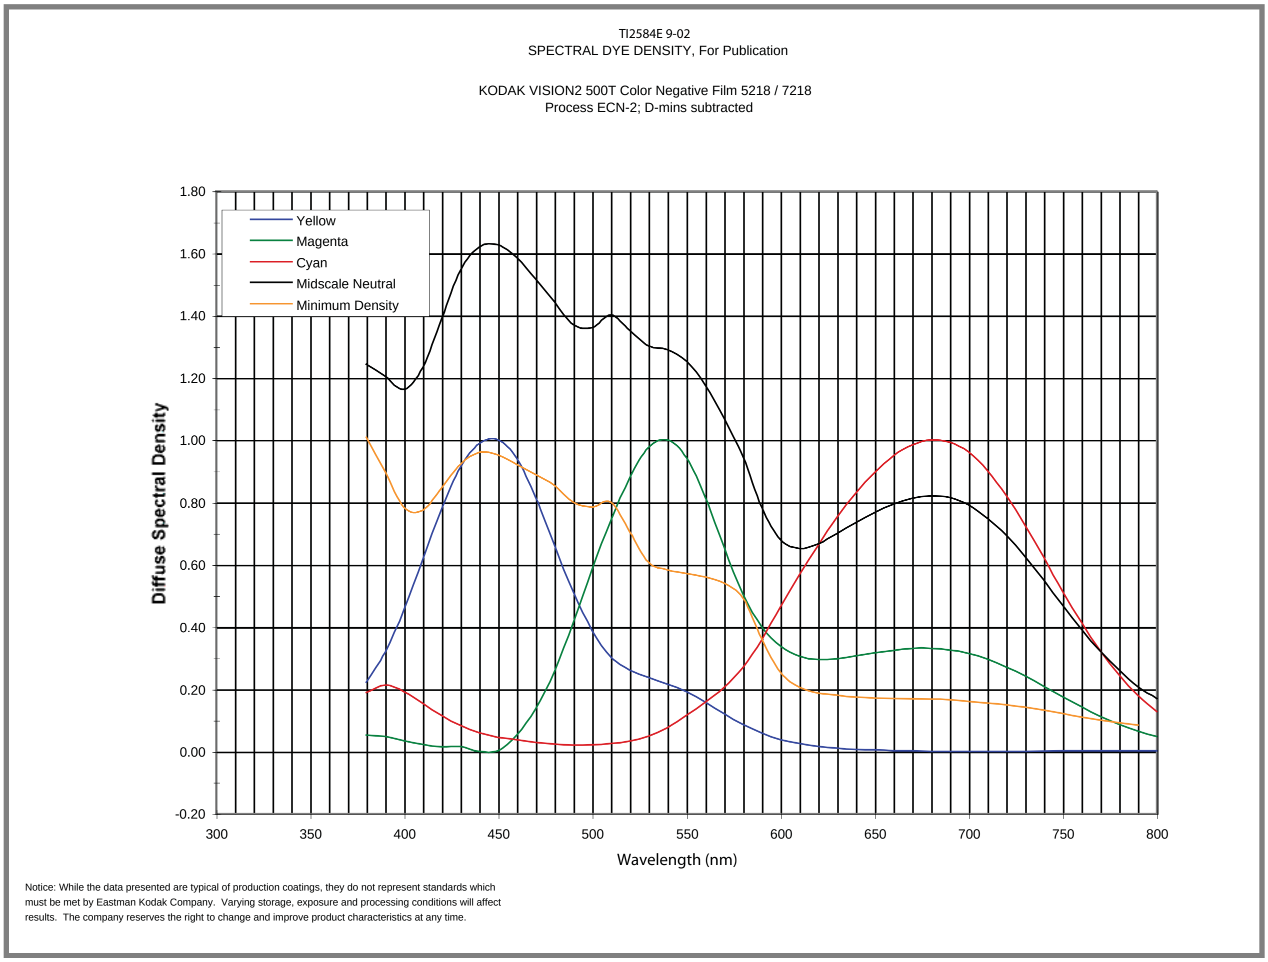

Most illuminants also have different power at different wavelengths and an illuminant can be represented by a spectral power distribution curve.

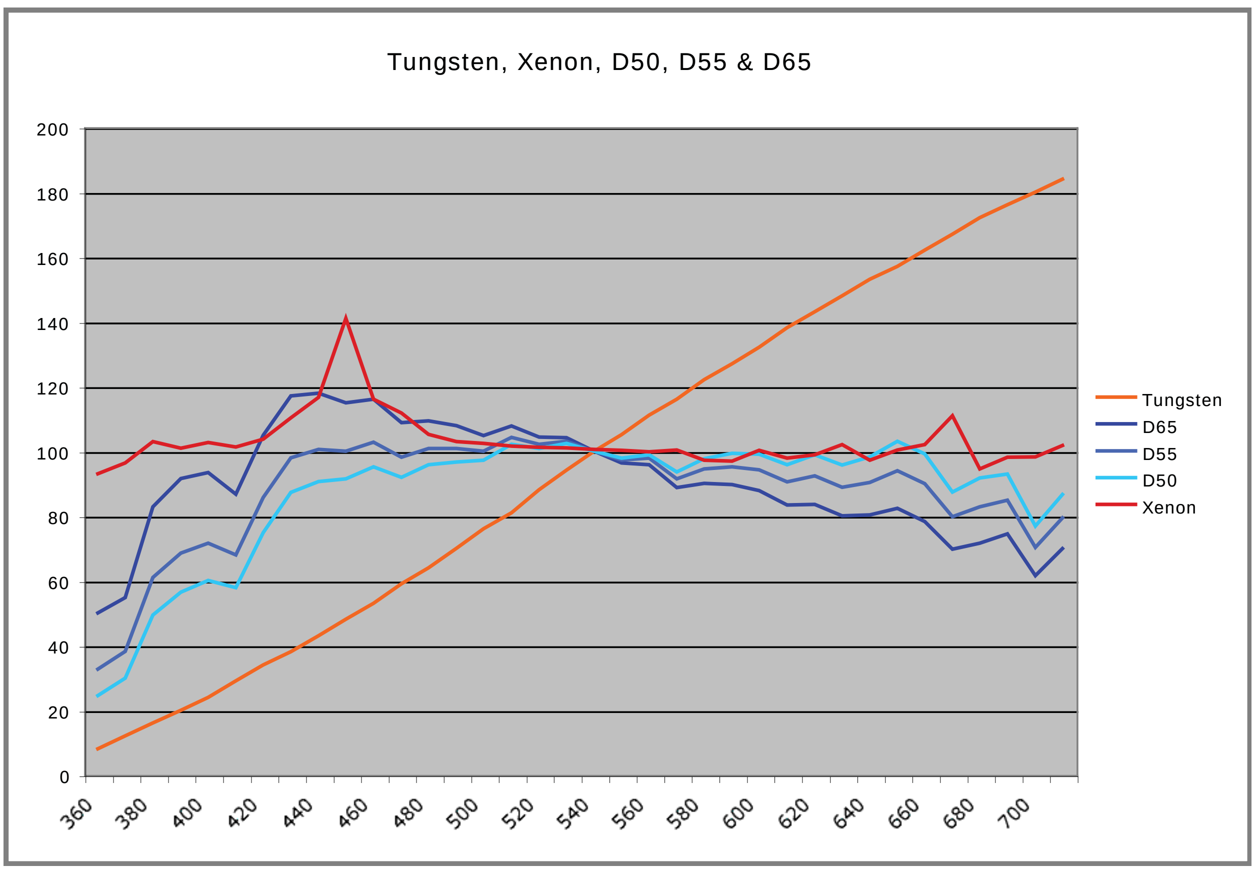

Finally most sensors respond differently to different wavelengths of light. Generally this response is further altered by the use of filters in front of the sensor. In the tri-linear array used in many scanners, each array has its own filter and its own spectral response which can be represented by a spectral sensitivity curve.

Kodak KLI Group 1 Sensor Spectral Sensitivity

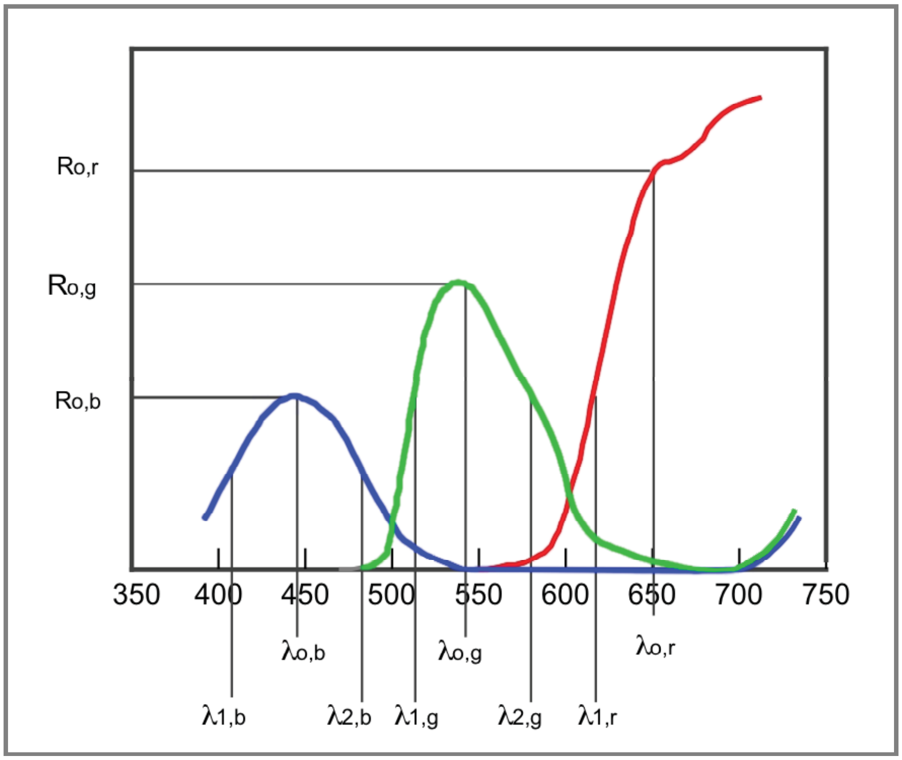

Design assumes an infrared filter in optical path.

To fully specify a type of density measurement, it is necessary to take into account the spectral power distribution of the illuminant and the spectral sensitivity curve of the sensor.

For any spectrally-selective object, such as a photographic imaging dye, there is no meaningful density measurement that yields some kind of generic or overall density representing the relationship of the "amount" of light transmitted to the "amount" of incident light. A density measurement is obtained by integrating the density at all the relevant wavelengths; but to be meaningful, this density measurement must be accompanied at least implicitly by a specification of this spectral data.

Density Measurement

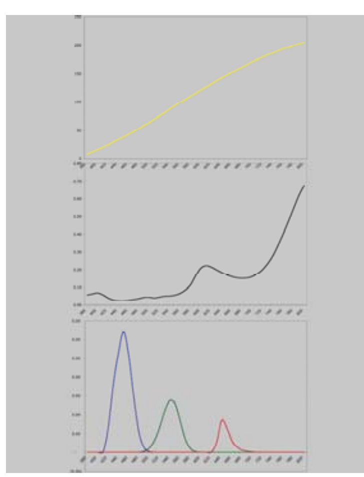

1) Illuminant Spectral Power Distribution
2) Material Linear Spectral Transmittance
3) Spectral Response of Sensor
   - Including spectral characteristics of complete optical path

Density Measurement

$$\text{Red Density} = \text{Log10}(Cr) \div \text{Log10}(Rr)$$

Where

$$Cr = \sum_{\lambda} I(\lambda) Sr(\lambda)$$

$$Rr = \sum_{\lambda} I(\lambda) T(\lambda) Sr(\lambda)$$

λ = wavelength, I = Illuminant,
T = Transmittance of Dye,
Sr = red sensitivity of sensor

The net response at each wavelength is the product of the linear value of each of these curves at that wavelength. In order to compute this response, the value for the dye density curve must be converted to a linear transmittance value. Sometimes other curves are also given as log curves, in which case they would also have to be converted to linear values for the calculations.

The standards for Status M and Status A densitometry include a spectral response curve that specifies the product of the spectral power distribution of the illuminant and the spectral sensitivity curve of a sensor for each color. For transmission measurements it represents the net spectral response of an "open gate" or a reading of a perfectly clear material.

It is possible with a densitometer like an Xrite 310 to switch back and forth between Status M and Status A and to take both types of density readings of the same piece of film. The result will be different red, green and blue densities for each type of reading, and how great the differences will be will depends on the actual color being measured. The reason for these differences is the difference in the spectral response specification for each type of density measurement.

Status A = 1.10, 0.97, 1.11
Status M = 1.24, 1.00, 1.16

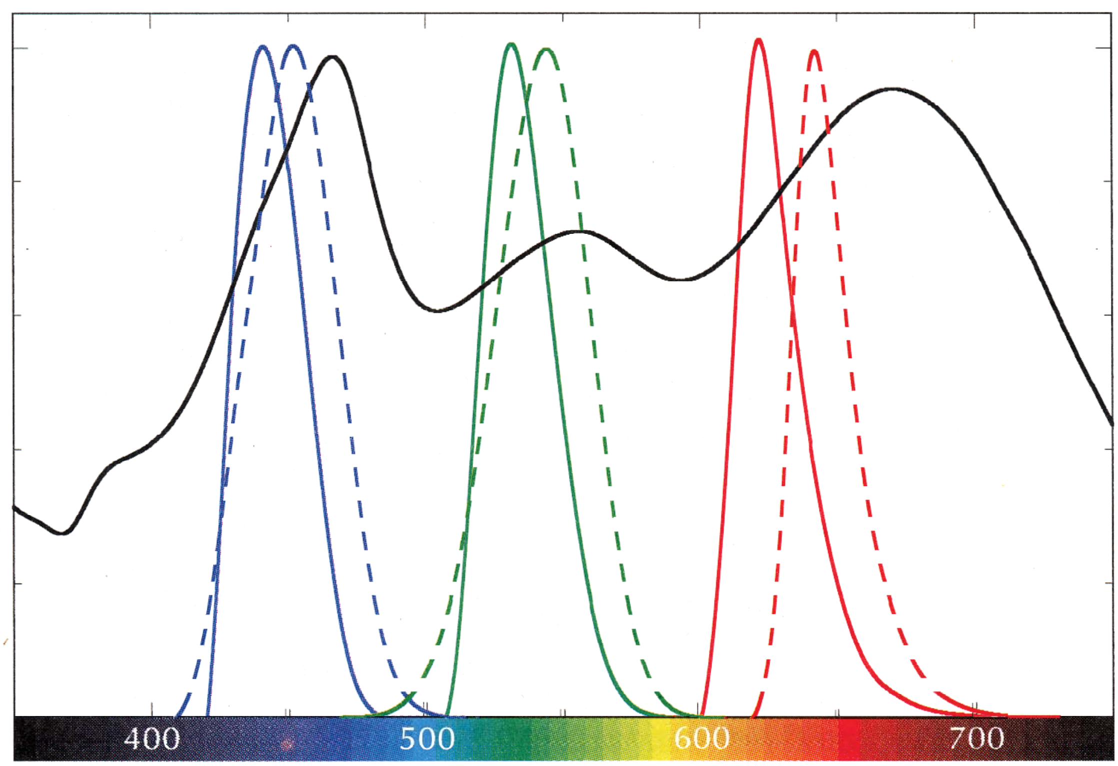

Wavelength, nm

Based on illustrations on pages 452 & 453 of Giorgianni's book

Status M Response and 5218 Dye Densities

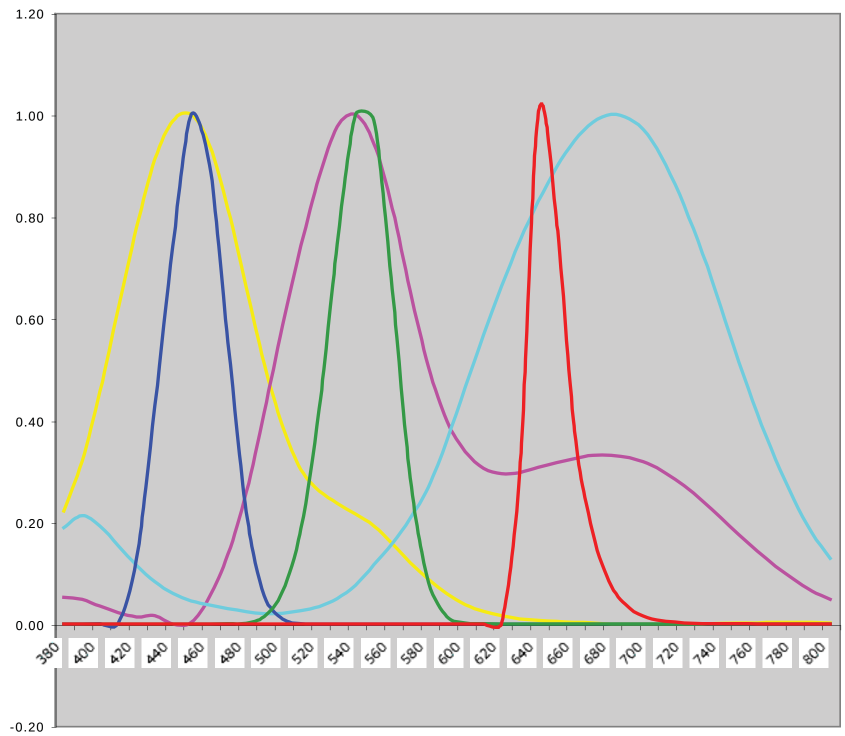

It is often assumed that since Status M is universally used for evaluating the results in processing negative, Status M would the natural choice for a scanner metric. The Imagica scanner may have been designed to yield Status M densities, but the original Kodak Cineon scanners as well as the scanners used in Kodak digital photofinishing operations have printing density responsivities; and it is important to understand why.

Some older explanations of Status M give the impression that it is also called printing density, (See for example page 247 of the 4th edition of The Reproduction of Colour by R.W.G. Hunt.) According to Ed Giorgianni, however, Status M was never intended to correspond to printing density. Status M was intended to measure negative film dyes more or less on peak to yield something more like analytical densities and in order to provide layer-by-layer information that is most useful for monitoring film manufacturing and chemical process control.

2383 Printing Density

Combination of Print Stock Spectral Sensitivity with Printer Illuminant and Optics

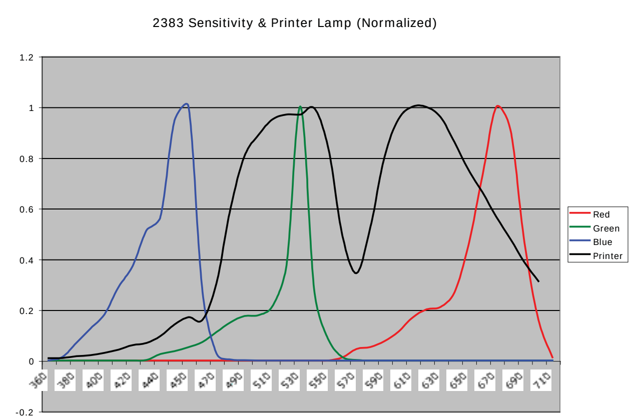

Printing density is generally described as the density of the negative as "seen" by the print stock. This means that the spectral sensitivity of the print stock combined with the spectral power distribution of the illuminant provided by the printer lamp house and optics define "printing density." As a result there is no single standardized universal printing density metric. Even though the illuminant in a printer can be standardized, the spectral response of print stock is subject to change as print stocks evolve or improve.

SMPTE Recommended Practice 180-1999 was an attempt to define a universal printing density. It was probably based on the Eastman color print stock commonly used around 1990, and everyone now seems to agree that it is obsolete. At any rate it clearly does not correspond to printing density for 2383.

At any given time, however, there are a limited number of prints stocks manufactured and each is designed for use with all the available camera negatives.

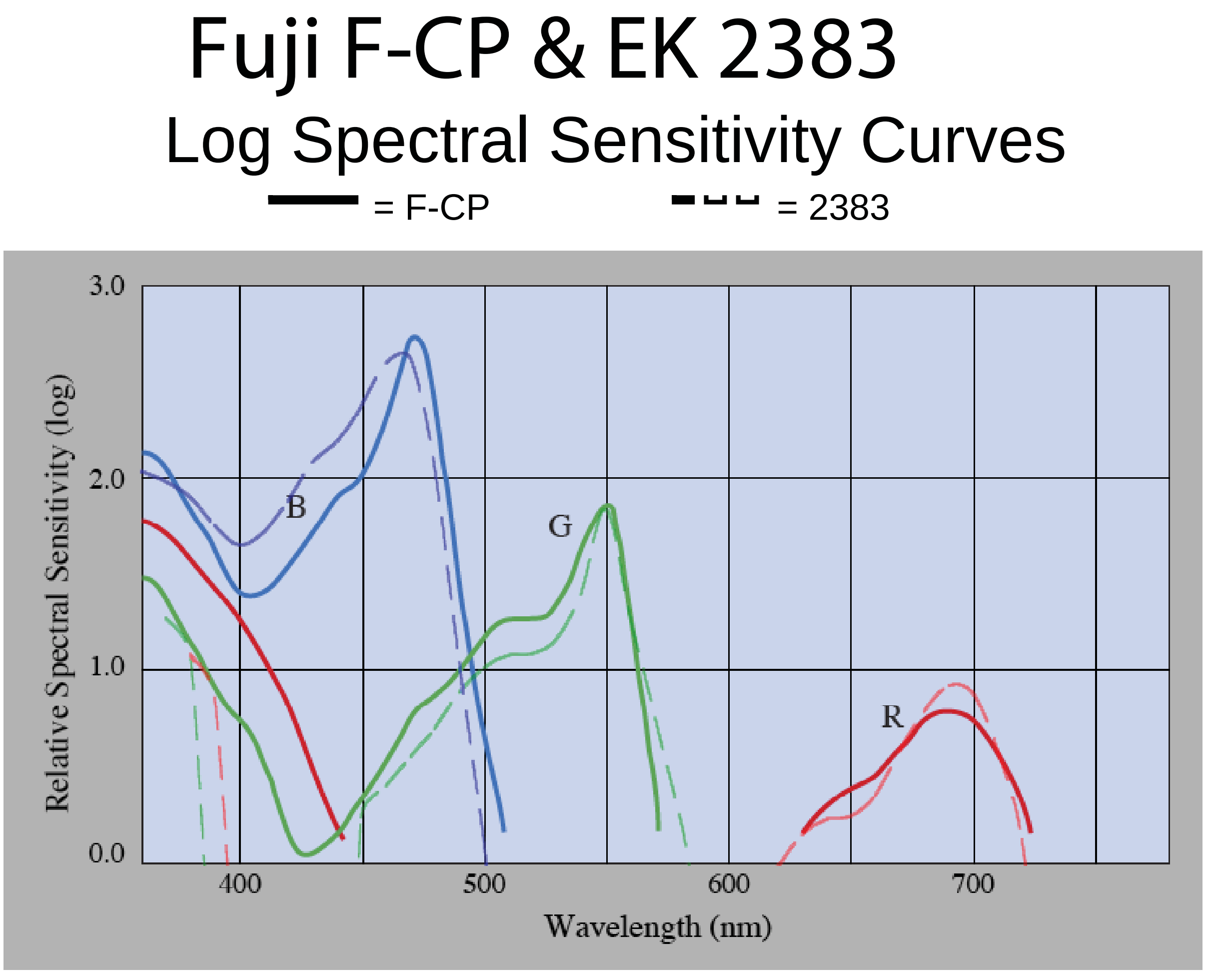

Currently the industry standard is undoubtedly Eastman Kodak 2383 and if we are going to talk about a printing density, it should be printing density for 2383. Some will argue that printing density should be defined relative to an intermediate stock, specifically Eastman Kodak 5242, since scanning is the electronic equivalent of duplicating the negative with an intermediate stock. For the moment we shall table this issue while noting that Eastman Kodak 5242 is designed to have a spectral sensitivity which emulates the spectral sensitivity of Kodak print stocks.

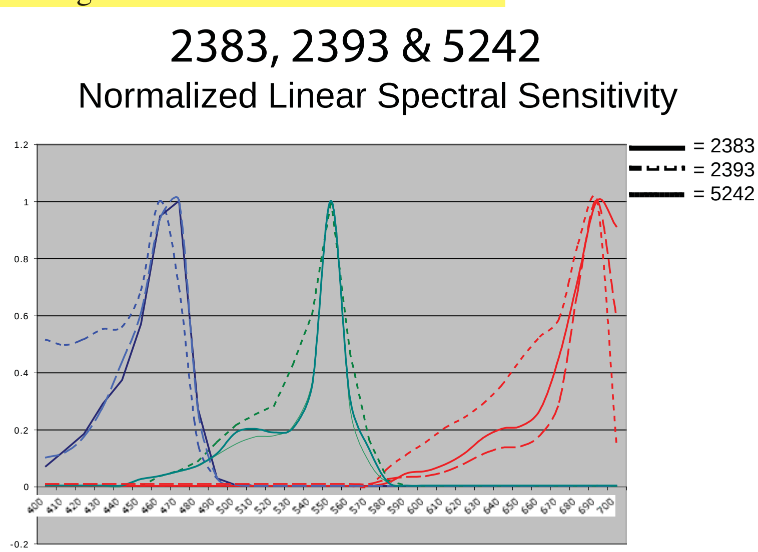

If we combine the spectral sensitivity curve for 2383 with the spectral power distribution curve for the lamp house and optical path in an industry standard printer, we have a spectral response curve that effectively defines printing density.

There is clearly a difference between Status M and printing density, but the question is how much difference this difference really makes in practical terms and whether it is practical to convert one to the other with sufficient accuracy.

2383 PD, Status M & RP 180

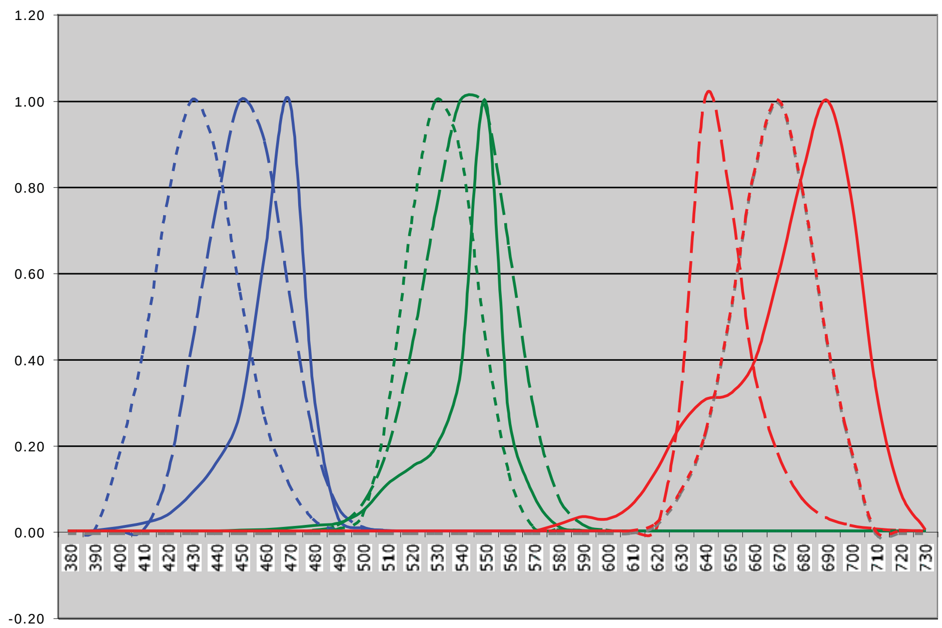

On a theoretical level the first thing that must be understood is that two negative stocks with different dye spectral density curves could result in the same Status M density reading but have different printing densities. Similarly a color patch on two different negatives will very commonly have the same printing density while having different Status M densities.

Printing Density & Status M for 5218 Neutral

= Relative Transmittance of Dye = Status M = 2383

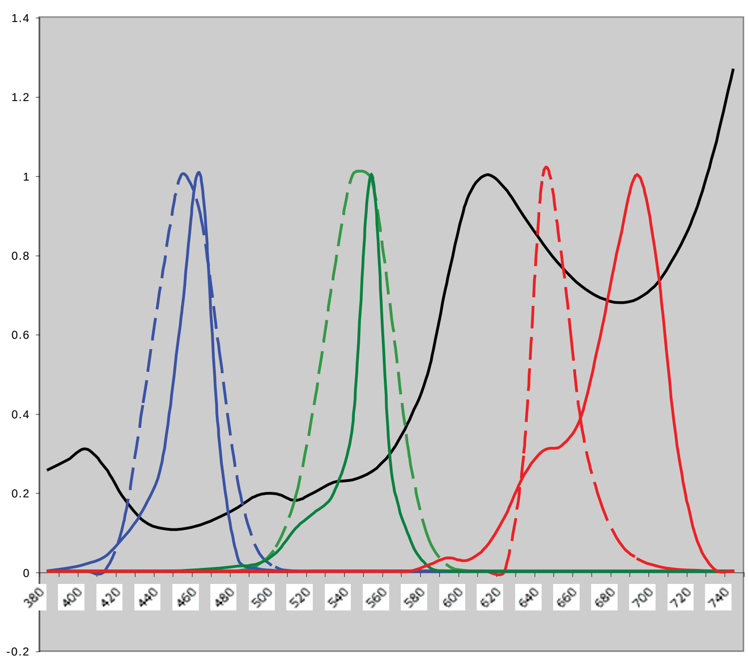

2383 Printing Density		
1.530	1.259	0.782
Status M		
1.565	1.245	0.755

The next thing we must acknowledge is that a scanner may well produce density values that are neither Status M nor printing density. If the spectral product of scanner illuminant power distribution and the scanner sensor spectral sensitivity differs sufficiently from the response curve for Status M or printing density, the raw output of the scanner will be something we should call "scanner density" since it is a measure of the density as "seen" by the scanner rather than as "seen" by 2383 print stock or a Status M densitometer.

5246 & 5218

Same Printing Density with Different Status M

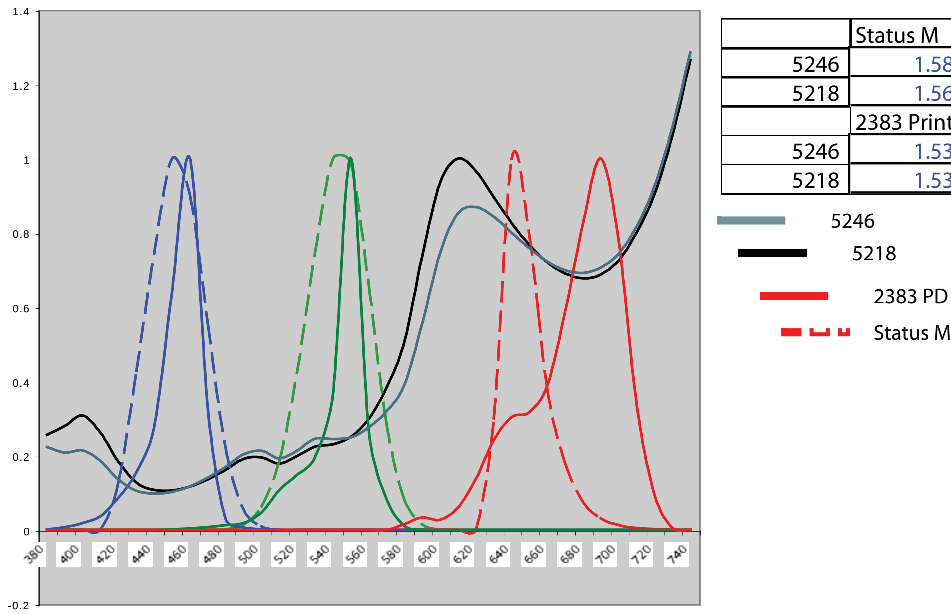

	Status M		
5246	1.583	1.248	0.764
5218	1.565	1.245	0.755
	2383 Printing Density		
5246	1.530	1.259	0.782
5218	1.530	1.259	0.782

Obviously in a perfect world we would design our scanner to match the responsivity of either Status M or printing density.

My impression is that the original Cineon scanner was designed to get as close to the printing density of the then current print stock as possible and that a 3x3 matrix had to be added to make it match more closely. This gives rise to the assumption that a mismatch in spectral response is no big deal since all that is required is a 3x3 matrix to correct for any error. I think we have to be very careful about making this assumption.

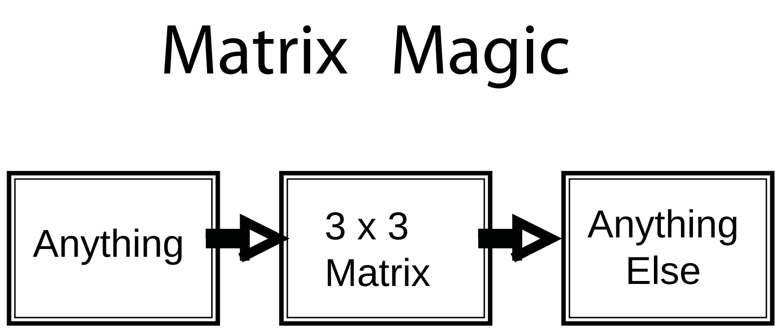

Because I myself have had to learn a lot of math during the last three years in my attempts to understand color science, I have concluded that many members of the committee have a basic misconception about what can be accomplished with 3x3 matrices and tend to indulge in a kind of hand wave of magical thinking associated with an invocation to simply "apply the appropriate matrix" to the data. I believe this is because 3x3 matrices are used in color science for two very different purposes, and it is easy to confuse them.

Many of the committee members are familiar with matrix algebra because of its use in 3D graphics for transformations in the orientation of an object in a three dimensional Cartesian coordinate system. This is an operation which involves three linear equations which can be accurately represented with a 3x3 matrix.

There is a corresponding use of a 3x3 matrix in color science when transformations are applied to convert an RGB color specification from a CIE color space based on one set of primaries to a CIE color space based on another set of primaries. The precision of this kind of operation is limited only by the precision of the floating point numbers employed.

3x3 Matrix Transforms in CIE Color Space

$$R_{709} = aR_{605} + bG_{605} + cB_{605}$$

$$G_{709} = dR_{605} + eG_{605} + fB_{605}$$

$$B_{709} = gR_{605} + hG_{605} + iB_{605}$$

$$\begin{vmatrix} R_{709} \\ G_{709} \\ B_{709} \end{vmatrix} = \begin{vmatrix} a\ b\ c \\ d\ e\ f \\ g\ h\ i \end{vmatrix} * \begin{vmatrix} R_{605} \\ G_{605} \\ B_{605} \end{vmatrix}$$

Other Uses of 3 x 3 Matrices in Color Science

e.g. Three Channel Camera Capture
To Scene-ReferredCIE XYZ Scene-Referred

$$X = aC_R + bC_G + cC_B$$

$$Y = dC_R + eC_G + fC_B$$

$$Z = gC_R + hC_G + iC_B$$

Matrix Transforms in 3D Graphics

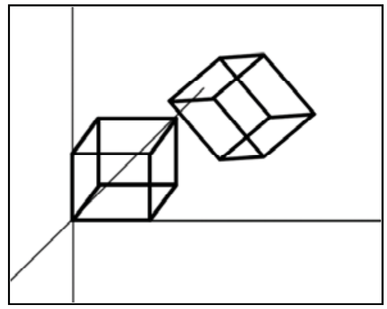

$x' = ax + by + cz + d$
$y' = ex + fy + gz + h$
$z' = ix + jy + kz + l$

There is, however, another use of 3x3 matrices in color science which does not enjoy this kind of precision. If you have a three dimensional color space which is defined by a set of spectral sensitivities such as the spectral sensitivities of the human visual system underlying the CIE color system and you want to convert a set of tristimulus values to a set of values in a color space defined by three different spectral sensitivity curves, you may no longer be dealing with three linear equations. In this case any 3x3 matrix you use is probably derived by numerical analysis from a specific set of samples of corresponding colors in both spaces. When the second set of spectral sensitivity curves is not a linear combination of the first, any 3x3 matrix produces only an approximation which is subject to varying degrees of error depending on where the color is in the color spaces.

Operations of this sort are such a fundamental ingredient in color science that a newcomer may just assume that such a use of a 3x3 matrix is the equivalent of the use of a matrix to do a transform in three dimensional geometry. It is not, and it is important to see that the spectral sensitivities of film, digital cameras and scanners or densitometers are not generally linear combinations of the color-matching functions of the CIE Standard Observer. This is, of course, the reason the procedures described by Jack are required to characterize a digital camera.

There is a difference between scanning a negative in which color is captured in essentially three dyes and capturing a scene in which color may be composed of any conceivable spectral power distribution. Dyes in a film emulsion follow Beer's law so that the densities at any two given wavelengths will always maintain the same ratio regardless of the total amount of dye. In other words the dye density curve can be scaled to represent varying dye amounts. This helps to simplify (for some) the math required to convert from one density metric to another and it means that there is a one-to-one relation-

Beer 's Law

Note: Applies to Density not Transmittance

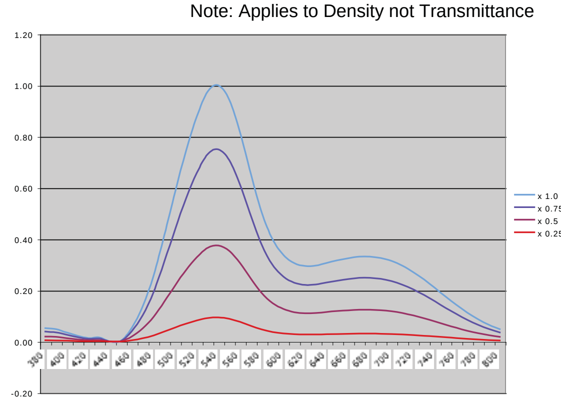

ship between two density metrics even if the relationship is not a linear one. For any given Status M density for a particular film there is only one corresponding printing density and vice versa.

So we have three different types of density that might be produced by a scanner: a unique scanner density, Status M density and printing density for a particular print stock. Each may have something to recommend it.

Encoding and exchanging the raw scanner density values would only be useful if there were also a profile for the scanner and for the negative stock which could be associated with the data. The only conceivable advantages offered by this approach might be that the data would be completely unambiguous and that the interpretation of the data from a scan could be improved as technology for characterizing the scanner improved. If the profile contains the spectral response curve for the scanner, an encoded scanner density value would be an unambiguous indication of the amount of dye present for each color component provided the spectral dye density information for the type of film scanned was known as well. (This is a little over-simplified because there are more than three dyes contained in a negative emulsion, but for illustrative purposes we are assuming only three "effective" dyes as represented by the dye density curves for the film.)

Knowing the actual dye amounts in the negative tells us nothing about the colors in the scene as photographed nor the color as it would appear in a print – unless we have a great deal more information about the negative stock and/or the print stock on which it would be printed. Since no one is seriously proposing that we encode raw scanner output values, this is an academic issue which can be pursued at someone's leisure.

Suppose we could design a scanner so that its spectral response matched exactly the spectral response defining printing density for Eastman Kodak 2383. In this case the raw scanner output would be printing density, and we would have an unambiguous indication of how the color will appear on a print using 2383. By empirical testing or modeling the print film and projection we can predict the CIE colorimetry of the color that will be projected on the screen. This will be true no matter what motion picture negative stock is scanned.

For output-referred colorimetry printing density is clearly the best metric to use for scanning. If the scanner's spectral response can not be made to match exactly the spectral response defining printing density for 2383, the closer it is the less error there will be in the prediction of the projected color. Processing derived through analysis of a number of color samples can be applied to the raw scanner output to minimize these errors, but there is no free lunch.

If scanner density must be converted to printing density because of a mismatch in spectral responsivity between the scanner and the print stock, the conversion required to minimize the errors will depend on the negative stock being scanned. If two negative stocks having different spectral dye curves are scanned, each will require a different conversion from scanner density to printing density. Unless a scanner can be designed so that its native spectral responsivity is the same as that of a print stock, proper scanning of motion picture negatives requires that the scanner operator adjust the scanner setup based on the type of negative stock being scanned. There is a very real difference between building printing density into the scanner and compensating for a different set of spectral responsivities with signal processing downstream from the scanner. This is, I confess, a complication which had escaped me until recently.

Obviously the same holds true for Status M, but even with a scanner whose spectral responsivities are an exact match to those defining Status M the conversion to printing density is stock dependent.

Suppose someone designed a scanner to have a spectral response matching exactly that for Status M. What would be gained by choosing this as the metric for encoding? We would again have an unambiguous indication of the amount of each dye contributing to each pixel in the frame (assuming

we have sufficient data about the negative stock that was scanned), but we would be one step further away from a prediction of the color as it would appear on a print. This is because there can be no universal conversion from Status M to printing density for all negative stocks.

The conversion from Status M to printing density involves the spectral dye density curves for the negative, and if the spectral dye density curves change (as they may from one negative stock to another) the relationship between the two density metrics changes. In other words in order to do the best output referred interpretation of Status M densities we need to know not only the dye density curve of the print but also the dye density curve of the negative. With printing density we do not need the dye density curve of the negative, and we do not even need to know what kind of negative it was in order to predict output colorimetry for a given print stock and projection.

For each particular negative stock there is a one-to-one relationship between Status M and printing density. From a practical point of view, however, the calculation of printing density from Status M density involves spectral samples with regression and is limited in its accuracy by the limits of the regression. (The regression is used to find a density factor or "amount of dye" that will yield the given Status M density. Then a simple spectral calculation using that density factor and the dye density curve can yield the corresponding printing density.) Doing these calculations for a large enough set of samples will permit a 3x3 matrix to be derived that can be used to approximate the spectral calculations. Note that this is still an approximation. The only way to get an accurate conversion is to use spectral calculations or to use polynomials instead of linear equations. According to Ed Giorgianni accurate conversions can generally be achieved with a polynomial in the form of

R'=a11R+ a12G+ a13B+ a14RxR+ a15RxG+ a16RxB+ a17

G'=a21R+ a22G+ a23B+ a24GxG+ a25GxR+ a26GxB+ a27

B'=a31R+ a32G+ a33B+ a34BxB+ a35BxR+ a36BxG+ a37

If one is primarily interested in interpreting negative density in terms of scene-referred exposure values, Status M may be a better choice for the scanner metric. This is because published data for the negative film stock generally includes a characteristic curve showing the relationship between Status M negative density and log exposure for a grayscale. However camera negative stocks are generally designed with interlayer interimage effects that complicate the calculation of channel independent RGB exposure from Status M or any other type of density measurement. There are also unwanted absorptions in the dyes that produce cross-talk in any form of integral density measurements. Nonetheless a three step conversion process can yield accurate results for a particular negative stock:

1) A procedure to convert density measurements into channel independent red, green and blue density values. As we have seen in the "unbuild" method from Tom and Gabe, this may need to be a polynomial or a 3D lookup table rather than just a 3x3 matrix.

2) The channel independent characteristic curve or a lookup table to derive log exposure from channel independent red, green and blue densities and

3) A 3x3 matrix to convert the RGB exposure to CIE XYZ or some other set of tristimulus values in the CIE color space

Note that the implementation of this conversion process is stock dependent. The conversion for 5218 will be different from the conversion for 5229.

It is worth noting at this juncture that if we derive printing densities from the scanner we can convert these to scene-referred values for a "universal generic" negative stock without knowing what the actual negative stock was. To the extent that the actual negative stock differed from the imaging characteristics ascribed to this generic universal stock, the scene-referred image will retain some of the characteristics of actual negative. This generic unbuild, however, is only possible because printing density represents the common ground between all negative stocks.

Since we are not designing a scanner ,we do not get to choose what the spectral response of the scanner will be. We have to work with existing scanners, and it is unclear to me how much information we have about the spectral responses of the existing scanners. We could, of course, make a recommendation for the choice of spectral response for future scanners or updated versions of current scanners. In the absence of accurate data on each scanner, we must rely on a calibration process based on appropriate film samples.

We are proposing a color management scheme involving the exchange of color data in an input color encoding space that requires the color information to be specified in terms of CIE colorimetry. This implies that any scanner output that is encoded as some form of density will have to be converted to input-referred channel-independent exposure values for encoding for exchange. If we provide calibrated film targets which can be used to generate transforms from scanner density to printing density along with the transform for converting printing density to input referred CIE colorimetry, the scanning facility can deliver either printing density values or CIE colorimetric input color encoding space values. The client will be free to request either, and the client will have from us the means to convert losslessly from one to the other at any point in the post-production workflow.

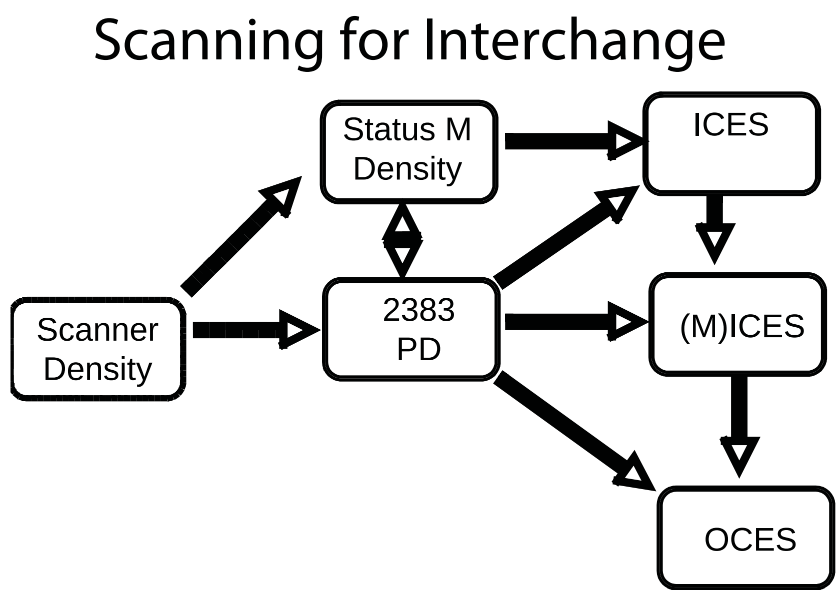

There is obviously some concern about maintaining and standardizing the current workflow based on densitometric encoding of film scans. If a density metric is going to be standardized for exchange of film scans, it should be printing density since the only reason to continue exchanging density values is to preserve the existing workflow with its emphasis on an output-referred interpretation of negative density.

When all is said and done, we may very well recommend procedures for scanning and exchanging images which do not result in the most accurate unbuilding of film or conversion from scanner density to printing density. If we do so, it is my hope that we shall do so with the full knowledge of where we have cut corners and that we shall present our recommendations in the context of a valid explanation of the ideal method for digitizing motion picture methods along with a clear demonstration of the magnitude of error which we have considered acceptable.

When I asked Ed Giorgianni for feedback on an early draft of this presentation, he responded with among other things the following points to bear in mind:

1) When multiple types of scanners are involved, scanner values alone are not interchangeable.
2) Except in the very unlikely case where a scanner's responsivities are a linear combination of some aim responsivities, such as Status M or Printing Density, transformations from scanner densities to aim densities will be film product dependent.
3) Accurate transformation from one type of integral densitometry to another almost always requires more than a simple 3x3 matrix.

4) Status M values are always film product dependent and cannot simply be used interchangeably.

5) Printing Density values are also film product dependent because the products themselves are different. However, Printing Density measurements represent those differences appropriately. Therefore, Printing Density values can be interchanged without further transformation if the desire is to maintain the original product differentiation.

6) A scanner having Printing Density responsivities measures Printing Density values directly, thus its values can be interchanged without further transformation if the desire is to maintain the original product differentiation.

7) Unbuilding can be done from almost any type of density measurement. However, in order to limit the number of unbuilding transformations needed to one per product, it would make sense to standardize scanner measurements. Status M might be the most convenient to implement, primarily because many installations have Status M densitometers that could be used to verify scanner measurements.

8) A single universal unbuilding transform can be used for all relevant negative films, if desired, only on a true Printing Density scanner. The decision as to whether to use a universal or product specific unbuilding transform would be based on the desired degree of product discrimination versus compatibility with other films and other forms of input.

9) A single universal unbuilding transform can be used, if desired, to convert Printing Density values computed on a non-Printing Density scanner. However, since the transformation from scanner to Printing Density is itself product dependent, the overall process still winds up being product specific. The decision as to whether to use a universal or product specific unbuilding transform then would be based on the desired degree of product discrimination versus compatibility with other films and other forms of input.

10) Meaningful unbuilding requires the use of a transformation from integral density to channel-independent density and the use of a set of channel-independent grayscale curves. Such transforms and curves are best provided by film manufactures.